

You Can Run, You Can Hide: The Epidemiology and Statistical Mechanics of Zombies

Alexander A. Alemi, Matthew Bierbaum, Christopher R. Myers, James P. Sethna
Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, NY 14853
(Dated: March 6, 2015)

We use a popular fictional disease, zombies, in order to introduce techniques used in modern epidemiology modelling, and ideas and techniques used in the numerical study of critical phenomena. We consider variants of zombie models, from fully connected continuous time dynamics to a full scale exact stochastic dynamic simulation of a zombie outbreak on the continental United States. Along the way, we offer a closed form analytical expression for the fully connected differential equation, and demonstrate that the single person per site two dimensional square lattice version of zombies lies in the percolation universality class. We end with a quantitative study of the full scale US outbreak, including the average susceptibility of different geographical regions.

I. INTRODUCTION

Zombies captivate the imagination. The idea of a deadly disease that not only kills its hosts, but turns those hosts into deadly vectors for the disease is scary enough to fuel an entire genre of horror stories and films. But at its root, zombism is just that: a (fictional) disease, and so should be amenable to same kind of analysis and study that more traditional diseases have long benefited from.

Much scholarly attention has focused on more traditional human diseases [10], but recently, academic attention has turned some amount of focus towards zombies as a unique and interesting modification of classic disease models. One of the first academic accounts of zombies was the 2009 article by Munz et al. [12], in which an early form of a compartmental model of zombism was introduced. Since then, there have been several interesting papers published including works that perform Bayesian estimations of the zombie disease parameters [22], look at how emotional factors impact the spread of zombies [16], using zombies to gain insight into models of politics [9], or the interaction of a zombie epidemic and social dynamics [11, 19]. Additional essays can be found in two books [4, 20], both collections of academic essays centered around zombism.

Besides the academic papers, zombies have seen a bit of a resurgence in fiction. Of particular note are the works of Max Brooks, including a very detailed *Zombie Survival Guide* [1], as well as an oral history of the first zombie war [2] in a hypothesized post outbreak world. In both these works Brooks discusses lots of details of zombies and their behavior often glossed over in other media. In particular, he makes the connection to disease explicit, describing zombies as the result of a hypothetical virus: *Solanum*.

Zombies form a wonderful model system to illustrate modern epidemiological tools drawn from statistical mechanics, computational chemistry, and mathematical modeling. It also forms an ideal vehicle for public outreach: the Center for Disease Control uses preparation for a zombie apocalypse [17, 18] to promote emergency

preparedness. In this work, we will build up to a full-scale simulation of a zombie outbreak in the continental United States, with realistic values drawn from the literature and popular culture (section V, simulation accessible online [14]). Before that, we shall use statistical mechanics to scrutinize the threshold of zombie virulence that determines whether humanity survives (section IV). Preceding that, we shall show how methods from computational chemistry can be used to simulate every individual heroic encounter between a human and a zombie (section III). But we begin by describing and analyzing a simple model of zombies (the *SZR* model) – the simplest and most natural generalization to the classic *SIR* (Susceptible-Infected-Recovered) model used to describe infectious disease spread in epidemiology.

II. SZR MODEL

We start with a simple model of zombies, the *SZR* model. There are three compartments in the model: *S* represents the susceptible population, in this case the uninfected humans; *Z* represents the infected state, in this case zombies; and *R* represents our removed state, in this case zombies that have been terminated by humans (canonically by destroying their brain so as to render them inoperable). There are two transitions possible: a human can become infected if they are bitten by a zombie, and a zombie can be destroyed by direct action by a human. There are two parameters governing these transitions: β , the bite parameter determines the probability by which a zombie will bite a human if they are in contact, and κ the kill parameter that gives the probability that a human kills the zombie. Rendered as a system of coupled differential equations, we obtain, for a particular interaction site:

$$\dot{S} = -\beta SZ \quad (1)$$

$$\dot{Z} = (\beta - \kappa)SZ \quad (2)$$

$$\dot{R} = \kappa SZ \quad (3)$$

Notice that these interactions are *density dependent*, in the sense that the rate at which we convert humans to

zombies and kill zombies is dependent on the total count of zombies and humans in this site. This is in contrast with most models of human diseases, which frequently adopt *frequency dependent* interactions wherein S, Z, R would be interpreted as the fraction of the population in that state.

This distinction will become stark once we consider large simulations with very inhomogeneous populations. By claiming that zombies can be modeled by a single bite parameter β that itself is a rate per person per unit time, we are claiming that a zombie in a block with 5,000 people would be one hundred times as effective at infecting new zombies as a zombie in a block with fifty people, similarly the zombie in question would be killed one hundred times faster. This would seem false for an ordinary disease like the flu, but in the case of zombies, we argue that it is appropriate, as zombies directly seek out hosts to infect, at which point the human and zombie engage in a dual to the (un)death.

To ease analysis we can nondimensionalize the equations by choosing a relevant population size N , and recasting in terms of the dimensionless time parameter $\tau = t\beta N$ and dimensionless virulence $\alpha = \kappa/\beta$

$$\begin{aligned}\frac{dS}{d\tau} &= -\frac{SZ}{N} \\ \frac{dZ}{d\tau} &= (1-\alpha)\frac{SZ}{N} \\ \frac{dR}{d\tau} &= \alpha\frac{SZ}{N}\end{aligned}\quad (4)$$

Unlike a traditional disease (e.g., as modeled by SIR), for the zombie model, we have a stable configuration when either the human or the zombie population is defeated ($S = 0$ or $Z = 0$). Furthermore, unlike SIR , SZ admits an analytical solution, assuming $R(0) = 0$, and with $Z_0 \equiv Z(0)$, $S_0 \equiv S(0)$:

$$P \equiv Z_0 + (1-\alpha)S_0 \quad (5)$$

$$\mu \equiv \frac{S_0}{Z_0}(1-\alpha) = \frac{P}{Z_0} - 1 \quad (6)$$

$$f(\tau) \equiv \frac{P\mu}{e^{\tau P/N} + \mu} \quad (7)$$

$$Z(\tau) = P - f(\tau) \quad (8)$$

$$S(\tau) = \frac{f(\tau)}{1-\alpha} \quad (9)$$

Given the analytical solution, it is clear to see that the sign of P governs whether we will eventually have humans or zombies in the final state. If $\alpha < 1$, $P > 0$, so

$$\lim_{\tau \rightarrow \infty} f(\tau) = 0 \quad (10)$$

$$\lim_{\tau \rightarrow \infty} Z(\tau) = P = Z_0 + (1-\alpha)S_0 \quad (11)$$

$$\lim_{\tau \rightarrow \infty} S(\tau) = 0 \quad (12)$$

and we will always flow to a final state composed of entirely zombies and no humans, where P denotes the number of zombies that survives.

If however, $\alpha > 1$, then humans are more effective at killing zombies than zombies are at biting humans, but if we start with enough zombies in the initial state, we can still convert all of the humans before they have time to kill all of the zombies.

In fact, we can recast the dynamics in terms of the variables $P \equiv Z + (1-\alpha)S$ and $\chi = S/Z$ to gain further insights. First note that

$$\frac{dP}{d\tau} = P' = Z' + (1-\alpha)S' \quad (13)$$

$$= (1-\alpha)\frac{SZ}{N} - (1-\alpha)\frac{SZ}{N} = 0 \quad (14)$$

So that P is a constant of the dynamics. As for χ

$$\chi' = \frac{S'}{Z} - \frac{SZ'}{Z^2} \quad (15)$$

$$= -\frac{S}{N} - (1-\alpha)\frac{S}{N}\frac{S}{Z} \quad (16)$$

$$= -\frac{S}{N}(1 + (1-\alpha))\chi \quad (17)$$

$$= -\frac{P}{N}\chi \quad (18)$$

So that if we choose $N = |P|$, we end up with the very simple dynamics:

$$P'(\tau) = 0 \quad (19)$$

$$P(\tau) = P_0 = Z(\tau) + (1-\alpha)S(\tau) = Z_0 + (1-\alpha)S_0 \quad (20)$$

$$\chi'(\tau) = \begin{cases} -\chi & P > 0 \\ +\chi & P < 0 \end{cases} \quad (21)$$

$$\chi(\tau) = \frac{S(\tau)}{Z(\tau)} = \chi_0 \begin{cases} e^{-\tau} & P > 0 \\ e^{+\tau} & P < 0 \end{cases} \quad (22)$$

$$\chi_0 \equiv \frac{S_0}{Z_0} \quad (23)$$

Here we see that the dynamics is simply an exponential decay or increase in the ratio of humans to zombies $\chi = S/Z$. The final populations in either case are easy to see due to the conservation of P . If zombies win we have

$$Z_\infty = Z_0 + (1-\alpha)S_0 \quad (24)$$

And if humans win

$$S_\infty = S_0 - \frac{Z_0}{\alpha - 1} \quad (25)$$

1. *SIR* model

This dynamics should be compared to the similarly nondimensionlized density dependent *SIR* model:

$$\frac{dS}{d\tau} = -\frac{SI}{N} \quad (26)$$

$$\frac{dI}{d\tau} = \left(\frac{S}{N} - \mu\right) I \quad (27)$$

$$\frac{dR}{d\tau} = \mu I \quad (28)$$

Where here $\tau = t\beta N$ as above, but $\mu = \nu/(\beta N) = R_0^{-1}$, because in the *SIR* model our infected population recovers on its own. This should be contrasted with *SZR*, where the process of infection and recovery have the same functional form, depending on the product SZ . This μ is the inverse of the usual R_0 parameter used to denote the infectivity of the *SIR*, here used to make a closer analogy to the *SZR* model. It is this parameter that principally governs whether we have an outbreak or not. Unlike the α parameter for *SZR* which depends only on our disease constants β, κ , the relevant virulence for the density dependent *SIR* model (μ) has a population dependence.

Notice that while the only stable configuration for the *SIR* model is when we have no infected population ($I = 0$), the *SZR* model is stable when either the humans or zombies are depleted ($S = 0$ or $Z = 0$).

Beyond that, the *SIR* model does not admit a closed form analytical solution, but we can find a parametric solution by dividing the first equation by the third, revealing.

$$S(\tau) = S_0 e^{-\frac{(R(\tau) - R_0)}{\mu N}} \quad (29)$$

And using the observation that in the limit of infinite time, no infected population can persist, we can choose N to be the total population

$$S_0 + I_0 + R_0 = N = S_\infty + R_\infty \quad (30)$$

and so obtain a transcendental equation for the recovered population at long times.

$$R_\infty = N - S_0 e^{-\frac{(R_\infty - R_0)}{\mu N}} \quad (31)$$

Unlike the *SZR* model, here we see that no matter how virulent the disease is, the epidemic will be self-limiting, and there will always have some susceptibles left at the end of the outbreak. This is a stark qualitative difference between zombies and more traditional *SIR* models, arising from the fact that the “recovery” of zombies is itself dependent on the presence of susceptibles.

To visually compare the difference, in Figure 1 we’ve shown example analytic dynamics for both *SIR* and *SZR*




FIG. 1. Example analytical dynamics for the *SIR* and *SZR* models with an initial population of 200 people, 199 uninfected and 1 infected. The (susceptible, infected, removed) population is shown in (blue, red, black). The *SZR* results are solid lines while the *SIR* results are lighter lines. For both models $\tau = t\beta N$ where N was taken to be the total population. For the *SZR* model α was chosen to be 0.6, while for the *SIR* model μ was chosen to be 0.6 to show similar evolutions. Notice that in this case, in *SZR* the human population disappears and we are left with zombies in the end, while the *SIR* model is self limiting, and only a fraction of the population ever becomes infected.





FIG. 2. Example Gillespie dynamics for the *SIR* and *SZR* models with the same parameter settings as Figure 1. The (susceptible, infected, removed) population is shown in (blue, red, black). The *SZR* results are solid lines while the *SIR* results are lighter lines. The two simulations were run with the same seed so as to match their dynamics at early times.

III. STOCHASTIC SIMULATION

While most previous studies modeling zombie population dynamics have been deterministic, things get more interesting when we try to model discrete populations. By treating the number of zombies and humans as continuous variables in the last section, we are ignoring the ran-

dom fluctuations that arise in small populations: even a ferociously virulent zombie infestation might fortuitously be killed early on by happy accident. Similar problems arise in chemical reactions: reactions involving two types of proteins in a cell can be described by chemical reaction kinetics evolving their concentrations (like our *SZR* equations 4), but if the number of such proteins is small, accurate predictions must simulate the individual binary reactions (each zombie battling each human). Interpreting our *SZR* transitions as reaction rates, gives us a system akin to a chemical reaction with two possible transitions:



When a human and zombie are in contact, the probability of a bite in a small period of time is given by the bite rate and the size of the populations of the two species ($\beta SZ dt$), and similarly for the probability of a kill. In order to efficiently simulate this dynamics, we use Gillespie dynamics [7], which efficiently uses the computer to sequentially calculate the result of each one-on-one battle.

The stochasticity gives more character to the simulation. The fully connected continuous dynamics modelled by the differential equation is straight forward: either the humans win and kill all of the zombies, or the zombies win and bite all of the humans. While the continuous approximation may be appropriate at intermediate stages of the infection where the total population is large and there are a non-trivial number of infected individuals, we will eventually be interested in simulating an actual outbreak on an inhomogeneous population lattice, where every new site will start with a single infected individual. But even though we may be interested in modeling the outbreak case ($\alpha < 1$), we would like to allow the possibility that the humans manage to defeat the outbreak before it really takes off. The stochastic Gillespie dynamics allows for this possibility.

In Figure 2 we've shown an example of a single stochastic simulation using the same parameter settings as those used in Figure 1. The stochastic trajectory overall tracks the analytic result, but at points in the simulation there may be more or less zombies than anticipated if the dice fall that way.

Another implication of stochastic dynamics is that it is not always guaranteed that an $\alpha < 0$ outbreak will take over the entire susceptible population. For the parameter settings used in Figure 1 and 2, namely $\alpha = 0.6$ with a population of 200 and one infected individual to start, the zombies win only 40% of the time. Additionally, the number of zombies we end with isn't fixed; as shown in Figure 3.

In fact, we can solve for the probability that an $\alpha < 1$ simulation will go extinct in the limit of large populations. We are interested in P_{ext} , the probability that the




FIG. 3. Distribution for final zombies over 100,000 Gillespie runs of the same settings as Figure 2. Not pictured are the 60% of runs that end with no zombies in the final state. Compare these to the analytical result, in which the final population of zombies would be 81.




FIG. 4. The observed fraction of simulations that end in an extinction for the zombie outbreak, for 1,000 runs of 10^4 individuals at various values of α (eqn. 33). The observed extinction probabilities agree with the expectation that they should go as α , here shown as the dashed blue line. This is the same behavior as the *SIR* model.

outbreak goes extinct. At the very beginning of the simulation, there is only one zombie, who will be killed with probability $\kappa/(\beta + \kappa)$. If we kill the first zombie before he bites anyone, we guarantee extinction. Otherwise, the zombie will bite another human, at which point we will have two independent zombie lines that we need to each cause to go extinct, which will occur with probability P_{ext}^2 . This allows us to solve:

$$P_{\text{ext}} = \frac{\kappa}{\beta + \kappa} + \frac{\beta}{\beta + \kappa} P_{\text{ext}}^2 \quad (32)$$

$$P_{\text{ext}} = \frac{\kappa}{\beta} = \alpha. \quad (33)$$

The probability of extinction is just given by our dimensionless inverse virulence α . In Figure 4 we've shown the

observed extinction probabilities for 1,000 Gillespie runs of a population of 10^4 individuals at various values of α , and overlaid our expected dependence of α .

This same extinction probability ($P_{\text{ext}} = \mu = R_0^{-1}$) is observed for the *SIR* model [10]. This is not a coincidence. In fact, in precisely the limit that is important for studying the probability of an extinction event, namely at early times with very large populations, the *SZR* model and *SIR* are effectively the same, since the population of susceptibles (S) is nearly constant. Writing S as $S_0 - \delta S$, we have:

$$\frac{dZ}{d\tau} = (1 - \alpha) \frac{S_0 Z}{N} - (1 - \alpha) \frac{\delta S Z}{N} \quad (34)$$

$$\frac{dI}{d\tau} = \left(1 - \frac{\mu N}{S_0}\right) \frac{S_0 I}{N} - (\mu N + \delta S) \frac{I}{N}. \quad (35)$$

Here as $\delta S \rightarrow 0$, the two models are the same with $\alpha = \mu N / S_0$, another indication that the density dependent *SIR* model's virulence is dependent on population size.

To get a better sense of the effect of the stochasticity, we can look at the mean fractional population in each state for various settings of α and choices for initial population size. The results are shown in Figure 5.




FIG. 5. Results from many Gillespie runs. One thousand different simulations are run for each cell. Each simulation starts with a single zombie or infected individual. The runs are run until they naturally terminate, either because the susceptible population is deleted, the zombie population is gone, or there are no more infected individuals. Each cell is colored according to the mean fraction of the population occurring in each state. The top row is for *SZR* simulations and the bottom row is for *SIR* simulations. In both cases N is chosen to be 100. Here the stark contrast between density dependent *SZR* and *SIR* is made apparent. Notice that density-dependent *SIR* is very strongly population dependent.

Shown are the fractional populations in the final state left for both the *SZR* model (top row) and *SIR* model (bottom row) for different parameter combinations of α and the initial population. In all cases, the N parameter was chosen to be 100. For each pixel, 1,000 independently seeded runs of Gillespie dynamics were calculated until completion. Looking at the *SZR* results in the top row, we can see that the dynamics is fairly independent of population size once the population size gets above around 100 individuals. The population dependence for lower population sizes is an effect of the stochasticity. We can clearly see a transition in the susceptible population near $\alpha = 1$ corresponding to where our continuous dynamics would show a sharp boundary. Here the boundary is blurred, again due to the stochasticity. The final dead zombie population R remains small for all values of α ; for extremely virulent zombies $\alpha \ll 1$, very few will be killed by the humans before all of the humans are converted, while in the other extreme few zombies are created so there are few to be killed.

Contrast these results with the density dependent *SIR* dynamics shown in the second row. There can be no infected individuals left in the end, so only the fraction of S and R in the final state are shown. The two transitions in *SIR* couple differently to the population of infected and susceptible. While our nondimensionalized *SZR* model has $Z' = (1 - \alpha)SZ/N$, our nondimensionalized *SIR* has $I' = (S/N - \mu)I$. This creates a very strong population dependence. The transition observed in the S population is largely independent of μ , except on the very small end. When we move to inhomogeneous population lattices this means that for the density dependent *SIR* model, the most important parameter governing whether a particular site has a breakout infection is the population of that site on the lattice.

IV. CRITICAL BEHAVIOR OF LATTICE MODEL

Until now, we've considered fully connected populations, where any infected individual can infect any susceptible individual. But surely, a zombie in New York cannot bite someone in Los Angeles. Studies of the spatial spread of infectious diseases is one of the applications of *network science*; social diseases spread among intimate contacts, Ebola spreads by personal contact in a network of caregivers, influenza can be spread by direct contact, through the air or by hand-to-mouth, hand-to-eye or hand-to-nose contact after exposure to a contaminated surface. For most diseases, 'long bonds' dominate the propagation to distant sites [13] – airplane flights take Ebola to new continents. Zombies do not fly airplanes, so our model is closer in spirit to the spread of certain agricultural infestations, where the disease spreads across a lattice of sites along the two-dimensional surface of the Earth (although not those in which pathogens are transported long distances by atmospheric currents).

To begin, we will consider a two dimensional lattice, where each site contains a single individual. Each individual is allowed to be in one of three states: S , Z , or R . The infection spreads through nearest neighbor bonds only. That is, a zombie can bite or be killed by any susceptible individuals in each of the four touching sites.

To make direct contact with our zombie model, the rate at which a susceptible cell is bitten is given by βZ where Z is the number of zombie neighbors (since S is one), and the rate at which a zombie site is killed is κS where S is the number of susceptible neighbors.

Because all state transitions in the SZR model depend only on Z - S contacts, for computational efficiency, we need only maintain a queue of all Z - S bonds, that is connections along which a human and zombie can interact. At each step of the simulation, one of these Z - S bonds is chosen at random, and with probability $\beta/(\beta + \kappa) = 1/(1 + \alpha)$, we bite the human, marking it as a zombie and visiting each of its neighbors. If any of its neighbors are human, we add that link to our queue. With probability $\kappa/(\beta + \kappa) = \alpha/(1 + \alpha)$ we kill the zombie, removing any of its links to neighboring humans from the queue. This process matches the stochastic dynamics of our zombie model operating on the lattice.

Simulating zombie outbreaks on fixed lattices, there is qualitatively different behavior for small α and large α . When α is large, the zombies do not spread very far, always being defeated by their neighboring humans. When α is very small, the zombies seem to grow until they infect the entire lattice. This suggests evidence of a phase transition. Technically, the presence of a phase transition would mean that if we could simulate our model on an infinite lattice, there should be some critical α (α_c), above which any outbreak will necessarily terminate. Below the critical value, we have the possibility (assuming we don't go extinct) of having the infection grow without bound, infecting a finite fraction of individuals, even on the infinite lattice. The SIR model has been demonstrated to undergo such a phase transition, and we expect the zombie model does as well.

The study of *critical phenomena* includes a series of techniques and analyses that enable us to study the property of these hypothetical phase transitions even on finite lattices. A major theme of critical phase transitions is that with the order parameter (the parameter governing the transition, in this case α) set to precisely the critical value, models show scale free behavior, meaning there is no natural length scale to the dynamics, and various physical parameters all are governed by power laws.

With α chosen to be precisely at the critical value, we expect to see fractal like growth (Fig. 6). Note that there are holes (surviving pockets of humans) of all sizes in the figure. This reflects the proximity to the threshold: the battle between zombies and humans is so evenly matched, that one gets an *emergent scale invariance* in the survival patterns. This is in keeping with studies of the SIR model, which shows a similar critical behavior and phase

transition [8].




FIG. 6. Example cluster resulting from the single population per site square lattice zombie model with periodic boundary conditions near the critical point $\alpha_c = 0.437344654(21)$ on a lattice of size 2048×2048 .

Systems near critical points with this kind of scale invariance fall into *universality classes*. Different systems (say, a real disease outbreak and a simple computational model) can in many ways act precisely the same on large scales near their transitions (allowing us to predict behavior without knowing the details of zombie-human (anti)social interactions). The SIR model on a two-dimensional lattice with a single person per site falls into the percolation universality class [3], though details of its cluster growth can differ [21]. Given that the SZR model has two second order couplings, it is of interest whether it falls into the same percolation universality class.

To extract the scaling behavior of our zombie infestation, we study the distribution $P(s, \alpha)$, the probability that a single zombie will generate an outbreak of size s at inverse virulence α . (An outbreak will be a fractal cluster in two dimensions, with ragged boundaries if it dies out before reaching the entire world.) At $\alpha = \alpha_c$ where the zombies and humans are equally matched, we have an emergent scale invariance. A large outbreak will appear to almost stop several times – it can be viewed as a sequence of medium-sized outbreaks triggering one another. Medium-sized outbreaks are composed of small outbreaks, which are in turn composed of tiny outbreaks. At threshold, each of these scales (large, medium, small) is related to the lower scale (medium, small, tiny) in the same fashion. Let us oversimplify to say that at criticality an outbreak of size $3s$ is formed by what would have

been three smaller outbreaks of size s which happened to trigger one another, and these in turn are formed by what would have been three outbreaks of size $s/3$. If the probabilities and form of this mutual triggering is the same at each scale, then it would not surprise us that many properties of the outbreaks would be the same, after rescaling the sizes by a factor of three. In particular, we expect at the critical point to find the probabilities of avalanches of size s to be related to the probabilities at size $s/3$ by some factor f :

$$P(s, \alpha_c) = fP(s/3, \alpha_c). \quad (36)$$

This formula implies that $P(s, \alpha_c) \propto s^{-\tau}$, with $\tau = \log(1/f)/\log(3)$. The distribution of epidemic infection rates is a power law.

Figure 7 shows a thorough test of this dependence for our zombie model, following a procedure akin to that of reference [21]. We simulated a zombie outbreak on a two-dimensional lattice with periodic boundary conditions starting with a single zombie. With the outbreak sizes following a power law distribution, the probability that a site belongs to a cluster of size n_s is $P_s = sn_s$, so that at the critical point $P_s \sim s^{1-\tau}$. Integrating from s to ∞ , the probability that a point belongs to a cluster of at least s in size $P_{\geq s}$ should at the critical point itself follow a powerlaw: $P_{\geq s} \sim s^{2-\tau}$. To find our critical point α_c , we ran many simulations until our integrated cluster size distribution followed a power law, using the interpolation methods of reference [21] to get a precise estimate of the critical point.

For zombies on a two dimensional lattice, this critical point occurs at $\alpha_c = 0.437344654(21)$, the resulting integrated cluster size distribution is shown at the top of Fig. 7. Percolation theory predicts $\tau = 187/91$ in two dimensions, and we test that prediction in the bottom part of Fig. 7. Here, if we were precisely at the critical point and the *SZR* model was in the percolation universality class, we would have a perfectly straight line. Notice the small scale our experimental results vary over several order of magnitude. The clear agreement convincingly shows that the zombie model on the two dimensional lattice is in the percolation universality class.

As an additional check, we computed the fractal dimension of our clusters near the critical point using box counting, a distribution for which is shown in Figure 8. We find a fractal dimension $D = 1.8946(14)$, compared to the exact percolation value of $D = 91/48 = 1.895833$.

Why did we need such an exhaustive test (many decades of scaling, many digits in our estimate of α_c)? On the one hand, a much smaller simulation could have told us that there was emergent scale invariance and fractal behavior near the transition; one or two decades of scaling should be convincing. But it turns out that there are multiple different universality classes for this kind of invasion process, and their exponents τ and D are rather similar. And a small error in α_c can produce large shifts in the resulting fits for τ and D – demanding efficient programming and fast computers to achieve a definitive




FIG. 7. The cumulative distribution of epidemic sizes for the two dimensional zombie model near the critical virulence. The critical point found was $\alpha_c = 0.437344654(21)$. The top plot shows the probability of a site being in a cluster of at least s in size ($P_{\geq s}$). The fact that it forms a straight line on a log-log plot indicates that $P_{\geq s}$ is a power law, and the slope is $2 - \tau$. For comparison, the blue line shows the powerlaw corresponding to the percolation critical exponent: $\tau = 187/91$. The bottom plot shows the same data times $s^{\tau-2}$ using the exponent from percolation theory. The plot is very nearly flat suggesting the percolation exponent accurately describes the zombie model.

answer.

We conclude that the single person per site zombie infestation, near the critical virulence, will on long length scales develop spatial infestation patterns that are well described by two-dimensional percolation theory.

V. US SCALE SIMULATION OF ZOMBIE OUTBREAK

Having explored the general behavior of the zombie model analytically, stochastically and on homogeneous single person lattices, we are prepared to simulate a full scale zombie outbreak.




FIG. 8. A histogram of the observed fractal dimension of the zombie epidemic clusters as measured by box counting. These give a measured value of $D = 1.8946(14)$, consist with the exact percolation fractal dimension of $D = 91/48 = 1.895833$.

A. Inhomogeneous Population Lattice

We will attempt to simulate a zombie outbreak occurring in the United States. This will be similar to our lattice simulation, but with an inhomogeneous population lattice. We based our lattice on code available for creating a “dot map” based off the 2010 US Census data [15]. The 2010 Census released census block level data, detailing the location and population of 11,155,486 different blocks in the United States. To cast these blocks down to a square grid, we assigned each of the 306,675,005 reported individuals a random location inside their corresponding census block, then gridded the population into a 1500×900 grid based on latitude and longitude coordinates. The resulting population lattice can be seen in the top half of Figure 9. You will see the presence of many empty grids, especially throughout the western United States. This disconnects the east and west coasts in a clearly artificial pattern – our zombies in practice will gradually wander through the empty grid points. To add in lattice connectivity, we did six iterations of binary closing (an image processing technique) on the population lattice and added it to the original. The effect was to add a single person to many vacant sites, taking our total population up to 307,407,336. The resulting population map is shown in the bottom half of Figure 9. This grid size corresponds to roughly 3 km square boxes. The most populated grid site is downtown New York City, with 299,616 individuals. The mean population of the occupied grid sites is 420, the median population of an occupied site is 13.

B. Augmented Model

In order to more ‘realistically’ simulate a zombie outbreak, we made two additions to our simplified SZR model. The first was to add a latent state E (Exposed).

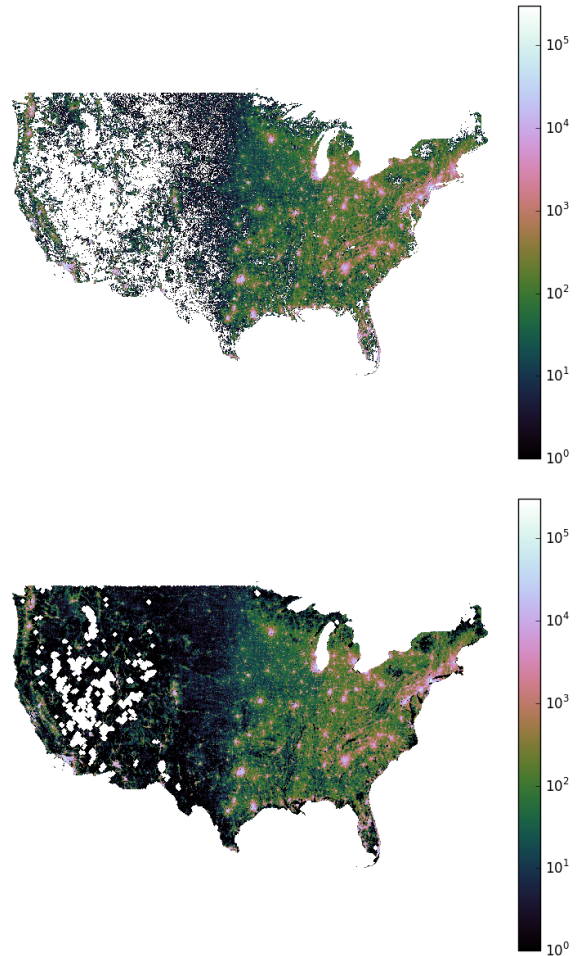


FIG. 9. A 1500×900 grid of the 2010 US Census Data. The above figure gives the raw results. Notice the multitude of squares with no people in them in the Western United States. The bottom figure shows the resulting map after 6 steps of binary closing added to the original population.

The second was to introduce motion for the zombies. Considered as a system of differential equations, we now have:

$$\dot{S}_i = -\beta S_i Z_i \quad (37)$$

$$\dot{E}_i = -\nu E_i \quad (38)$$

$$\dot{Z}_i = \nu E_i - \kappa S_i Z_i \quad (39)$$

$$\dot{R}_i = \kappa S_i Z_i \quad (40)$$

$$\dot{Z}_i = \mu \sum_{\langle j \rangle} Z_j - \mu Z_i \quad (41)$$

or as a set of reactions:

$$(S_i, E_i) \xrightarrow{\beta S_i Z_i} (S_i - 1, E_i + 1) \quad (42)$$

$$(Z_i, E_i) \xrightarrow{\nu E_i} (Z_i + 1, E_i - 1) \quad (43)$$

$$(Z_i, R_i) \xrightarrow{\kappa S_i Z_i} (Z_i - 1, R_i + 1) \quad (44)$$

$$\langle i j \rangle : (Z_i, Z_j) \xrightarrow{\mu Z_i} (Z_i - 1, Z_j + 1) . \quad (45)$$

Here i denotes a particular site on our lattice. $\langle j \rangle$ denotes a sum over nearest neighbor sites, $\langle i j \rangle$ denotes that i and j are nearest neighbors. In this model, zombies and humans only interact if they are at the same site, but the zombies diffuse on the lattice, being allowed to move to a neighboring site with probability proportional to their population and some diffusion constant (μ). We assume that the humans do not move, not only for computational efficiency, but because, as we will see, the zombie outbreaks tend to happen rather quickly, and we expect large transportation networks to shut down in the first days, pinning most people to their homes. The addition of a latent state coincides with the common depiction that once a human has been bitten, it typically takes some amount of time before they die and reanimate as a zombie. If a human is bitten, they transition to the E state, where at some constant rate (ν) they convert into the zombie state.

To choose our parameters we tried to reflect common depictions of zombies in movies. In the work of Witkowski and Blais [22], they performed a Bayesian fit of a very similar SZR model to two films, *Night of the Living Dead*, and *Shawn of the Dead*. In both cases, the observed α was very close to 0.8. This means that the zombies in the films are 1.25 times more effective at biting humans than the humans are at killing the zombies. We will adopt this value for our simulation. For our latent state, we adopt a value close to that reported for *Shawn of the Dead*, namely a half life of 30 minutes. To set our movement parameter, we estimate that zombies move at around 1 ft/sec. To estimate the rate at which the zombies will transition from one cell to the next, we assume that the zombies behave like a random gas inside the cell, so that the probability that a zombie will cross a cell boundary is roughly $\frac{1}{4} \frac{Z}{L^2} Lv \Delta t$, that is, one fourth of the zombies within $v \Delta t$ of the edge will move across that edge in a small amount of time. This suggests a value of μ of 0.0914 /hr. This corresponds to an average time between transitions of around 11 hours, which for a zombie stumbling around a 3 km block agrees with our intuitions. Finally, to set a rate for our bite parameter, we similarly assume that the zombies are undergoing random motion inside the cell at 1 ft/sec, and they interact with a human anytime they come within 100 feet. We can then estimate the rate at which humans and zombies will interact as $SZ \frac{Rv \Delta t}{L^2}$, which corresponds to a choice of β of around 3.6×10^{-3} /hr. Another way to make sense of these parameter choices is to ask how many susceptible individuals must be in a cell before a single zombie

β	3.6×10^{-3} /hr/person
α	0.8
κ	$\alpha\beta$
η	2 /hr
μ	0.0914 /hr

TABLE I. The parameters chosen for our US-scale simulations of a zombie outbreak. These parameters were chosen to correspond with standard depictions of zombies and simple physical estimations explained in the main text.

has a higher rate for biting a human than transitioning to a neighboring cell. For our choice of parameters, this gives

$$N\beta = 4\mu \implies N \sim 102 . \quad (46)$$

This corresponds to a low population density of ~ 11 people/km², again agreeing with our intuition. All of our parameter choices are summarized in Table I.

C. Simulation Details

To effectively simulate an outbreak at this scale, we employed the Next Reaction Method of [6]. We maintained a priority queue of all possible reactions, assigning each the time at which the reaction would take place, an exponentially distributed random number with scale set by the rate for the reaction. At each time step of the simulation, we popped the next reaction off of the queue, and updated the state of the relevant squares on our grid. Whenever population counts changed, we of course needed to update the times for the reactions that depend on those population counts. This method remained efficient for simulating the entire US. However, at late times a large amount of simulation time was spent simulating the diffusion of the zombies back and forth between highly populated states. We could have achieved additional computational efficiency by adopting the time dependent propensity function approach of Fu et al. [5].

D. Results

With the simulation in place, we are now in a position to simulate a full scale zombie outbreak. We first consider an outbreak that began with one in every million individuals starting in the Exposed (E) state in the United States. For a single instance the overall populations are shown in Figure 10. This looks similar to the analytical outbreaks we saw in Figure 1, but with a steeper rate of initial infection and some slight perturbations to the curves. The total population curves however hide most of the interesting features. In Figure 11 we attempt to give a sense of how this outbreak evolves, showing the state




FIG. 10. The S (blue), Z (red), R (black), and E (green) populations as a function of time for a full scale zombie outbreak in the continental United States starting with one in every million people infected.

of the United States at various times after the outbreak begins.

As you can see, for the parameters we chose, most of the United States population has been turned into zombies by the first week, while the geographic map doesn't necessarily seem all that compelling. In the early stages of the outbreak, while the population is roughly homogeneous, the zombie plague spreads out in roughly uniform circles, where the speed of the infection is tied to the local population density. Infestations on the coasts, with their higher population density, have spread farther than those near the center of the country. After several weeks, the map exhibits stronger anisotropy, as we spread over larger geographical areas and the zombie front is influenced by large inhomogeneities in population density. After four weeks, much of the United States has fallen, but it takes a very long time for the zombies to diffuse and capture the remaining portions of the United States. Even four months in, remote areas of Montana and Nevada remain zombie free.

To investigate the geographical characteristics of the outbreak, we must move beyond a single instance of an outbreak and study how different regions are affected in an ensemble of outbreaks. If it takes a month to develop and distribute an effective vaccine (or an effective strategy for zombie decapitation), what regions should one locate the zombie-fighting headquarters? We ran 7,000 different 28-day zombie outbreaks in the continental United States starting with a single individual. A single instance of one of these outbreaks originating in New York City is shown in Figure 12.

By averaging over all of these runs, we can start to build a zombie susceptibility map, as shown in Figure 13. In the top plot, we show the probability that the given cell is overrun by zombies after seven days. Here you can clearly see that there are certain regions – those surrounding populous metropolitan areas – that are at a greater risk. This is partly because those regions have lots of individuals who could potential serve as patient zero, and partly due to the rapid spread of zombies in




FIG. 11. Simulation of a zombie outbreak in the continental United States. Initially one in every million individuals was infected at random. Results are shown above at (a) one day, (b) two days, (c) one week, (d) two weeks, (e) three weeks, (f) four weeks, and (g) two months after the outbreak begins. Shown here are the population of susceptible individuals (S) in blue, scaled logarithmically, zombies in red and removed in green. All three channels are superimposed.

those areas. In the bottom plot, we plot the probability that the cell is overrun, but at the 28 day mark.

After 28 days, it is not the largest metropolitan areas that suffer the greatest risk, but the regions located between large metropolitan areas. For instance, in California it is the region near Bakersfield in the San Joaquin Valley that is at the greatest risk as this area will be overrun by zombies whether they originate in the San Francisco area or the Los Angeles / San Diego area. The area with the greatest one month zombie risk is north eastern Pennsylvania, itself being susceptible to outbreaks originating in any of the large metropolitan areas on the east coast.




FIG. 12. Status of the United States 28 days after an outbreak that started in New York City. Here blue represents humans, red represents zombies and green represents dead zombies. The three color channels have been laid on top of one another.

VI. CONCLUSION

Zombies offer a fun framework for introducing many modern concepts from epidemiology and critical phenom-

ena. We have described and analyzed various zombie models, from one describing deterministic dynamics in a well-mixed system to a full scale US epidemic. We have given a closed form analytical solution to the well-mixed dynamic differential equation model. We compared the stochastic dynamics to a comparable density-dependent *SIR* model. We investigated the critical behavior of the single person per site two-dimensional square lattice zombie model and demonstrated it is in the percolation universality class. We ran full scale simulations of a zombie epidemic, incorporating each human in the continental United States, and discussed the geographical implications for survival.

VII. ACKNOWLEDGEMENTS

We acknowledge NSF IIS-1247696 and Cornell University for support of this research, and thank Paul Ginsparg for useful references and conversations.

-
- [1] Max Brooks. *The Zombie Survival Guide: Complete protection from the living dead*. Broadway books, 2003.
 - [2] Max Brooks. *World War Z*. Querido's Uitgeverij BV, Em., 2013.
 - [3] John L Cardy and Peter Grassberger. Epidemic models and percolation. *Journal of Physics A: Mathematical and General*, 18(6):L267, 1985.
 - [4] Andrew Cartmel et al. *Mathematical Modelling of Zombies*. University of Ottawa Press, 2014.
 - [5] Jin Fu, Sheng Wu, Hong Li, and Linda R Petzold. The time dependent propensity function for acceleration of spatial stochastic simulation of reaction-diffusion systems. *Journal of Computational Physics*, 274:524–549, 2014.
 - [6] Michael A Gibson and Jehoshua Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The journal of physical chemistry A*, 104(9):1876–1889, 2000.
 - [7] Daniel T Gillespie, Andreas Hellander, and Linda R Petzold. Perspective: Stochastic algorithms for chemical kinetics. *The Journal of chemical physics*, 138(17):170901, 2013.
 - [8] Peter Grassberger. On the critical behavior of the general epidemic process and dynamical percolation. *Mathematical Biosciences*, 63(2):157–172, 1983.
 - [9] Ronald Hochreiter and Christoph Waldhauser. Zombie politics: Evolutionary algorithms to counteract the spread of negative opinions. *arXiv preprint arXiv:1401.6420*, 2014.
 - [10] Matt J Keeling and Pejman Rohani. *Modeling infectious diseases in humans and animals*. Princeton University Press, 2008.
 - [11] Bruno Mota. Optimum survival strategies against zombie infestations—a population dynamics approach. *Bulletin of the American Physical Society*, 59, 2014.
 - [12] Philip Munz, Ioan Hudea, Joe Imad, and Robert J Smith? When zombies attack!: mathematical modelling of an outbreak of zombie infection. *Infectious Disease Modelling Research Progress*, 4:133–150, 2009.
 - [13] Mark EJ Newman, I Jensen, and RM Ziff. Percolation and epidemics in a two-dimensional small world. *Physical Review E*, 65(2):021904, 2002.
 - [14] <http://matthierbaum.github.io/zombies-usa/>.
 - [15] <https://github.com/meetar/dotmap>.
 - [16] Felipe Nunez, Cesar Ravello, Hector Urbina, and Tomas Perez-Acle. A rule-based model of a hypothetical zombie outbreak: Insights on the role of emotional factors during behavioral adaptation of an artificial population. *arXiv preprint arXiv:1210.4469*, 2012.
 - [17] U.S. Department of Health, Human Services Centers for Disease Control, and Prevention. Preparedness 101: Zombie pandemic. <http://www.cdc.gov/phpr/zombies/>.
 - [18] U.S. Department of Health, Human Services Centers for Disease Control, and Prevention. Preparedness 101: Zombie apocalypse. <http://blogs.cdc.gov/publichealthmatters/2011/05/preparedness-101-zombie-apocalypse/>, 2011.
 - [19] Evelyn Sander and Chad M Topaz. The zombie swarm: Epidemics in the presence of social attraction and repulsion. *Mathematical Modelling of Zombies*, page 265, 2014.
 - [20] Robert Smith? *Braaaaaiiiiinnssss!: From Academics to Zombies*. University of Ottawa Press, 2011.

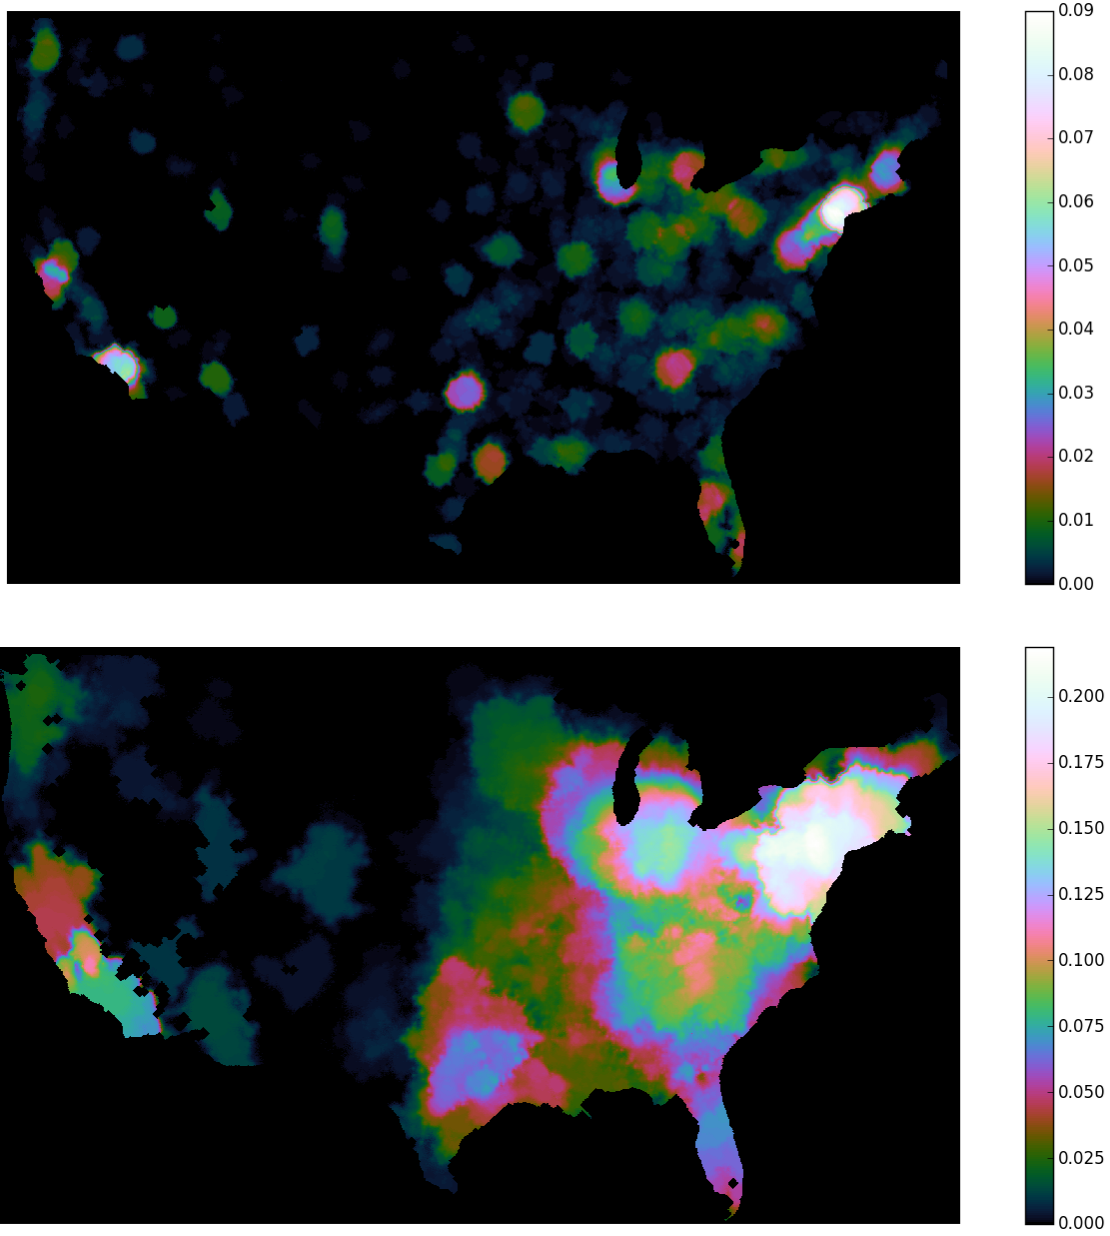


FIG. 13. Average survivability from US scale runs. In both cases, the plot shows the probability of being infected in that square after an epidemic that originates from a single infected individual chosen at random from the total population. The top figure is the probability of being infected after 7 days, while the bottom plot is after 28 days. In total, this represents 7,000 simulated runs starting from a single individual. The top plot represents the 1,467 outbreaks that lasted at least 7 days, the bottom plot represents 1,458 outbreaks that lasted at least 28 days.

[21] Tânia Tomé and Robert M Ziff. Critical behavior of the susceptible-infected-recovered model on a square lattice. *Physical Review E*, 82(5):051921, 2010.

[22] Caitlyn Witkowski and Brian Blais. Bayesian analysis of epidemics-zombies, influenza, and other diseases. *arXiv preprint arXiv:1311.6376*, 2013.